

Fast Rates by Transferring from Auxiliary Hypotheses

Ilja Kuzborskij

Idiap Research Institute
Rue Marconi 19, Martigny, Switzerland
ilja.kuzborskij@idiap.ch

Francesco Orabona

Yahoo! Labs
229 West 43rd Street, 10036 New York, NY, USA
francesco@orabona.com

July 20, 2015

Abstract

In this work we consider the learning setting where, in addition to the training set, the learner receives a collection of auxiliary hypotheses originating from other tasks. We focus on a broad class of ERM-based linear algorithms that can be instantiated with any non-negative smooth loss function and any strongly convex regularizer. We establish generalization and excess risk bounds, showing that, if the algorithm is fed with a good combination of source hypotheses, generalization happens at the fast rate $\mathcal{O}(1/m)$ instead of the usual $\mathcal{O}(1/\sqrt{m})$. On the other hand, if the source hypotheses combination is a misfit for the target task, we recover the usual learning rate. As a byproduct of our study, we also prove a new bound on the Rademacher complexity of the smooth loss class under weaker assumptions compared to previous works.

1 Introduction

In the standard supervised machine learning setting the learner receives a set of labeled examples, known as the training set. However, very often we have additional information at hand that could be beneficial to the learning process. One such example is the use of unlabeled data drawn from the marginal distributions, that gives rise to the semi-supervised learning setting [Chapelle et al., 2006]. Another example is when the training data is coming from a related problem, as in multi-task learning [Caruana, 1998], domain adaptation [Ben-David et al., 2010, Mansour et al., 2009], and transfer learning [Pan and Yang, 2010, Taylor and Stone, 2009]. Among others, there is the use of structural information, such as taxonomy, different views on the same data [Blum and Mitchell, 1998], or even a sort of privileged information [Vapnik and Vashist, 2009, Sharmanska et al., 2013]. In the recent years all these directions have received a considerable empirical and theoretical attention.

In this work we focus on a less theoretically studied direction in the use of supplementary information – learning with *auxiliary hypotheses*, that is classifiers or regressors originating from another tasks. In particular, in addition to the training set we assume that the learner is supplied with a collection of hypotheses and their predictions on the training set itself. The goal of the learner is to figure out which hypotheses are helpful and use them to improve the prediction performance of the trained classifier. We will call these auxiliary hypotheses the *source* hypotheses and we will say that helpful ones accelerate the learning on the *target* task. We focus on the linear setting, that is, we train a linear¹ classifier and the source hypotheses are used

¹Non-linear classifiers can be easily produced with the use of kernels.

additively in the prediction process, weighted by arbitrary weights. This generalizes the setting in which the outputs of the source hypotheses are concatenated with the feature vector, a widely used heuristic [Bergamo and Torresani, 2014, Li et al., 2010, Tommasi et al., 2014].

The scenario described above is related to the Transfer Learning (TL) and Domain Adaptation (DA) ones, or learning effectively from possibly small amount of data by reusing the prior knowledge [Thrun and Pratt, 1998, Pan and Yang, 2010, Taylor and Stone, 2009, Ben-David et al., 2010]. However, transferring from hypotheses offers an advantage compared to TL and DA frameworks, where one requires access to the data of the *source* domain. For example, in DA [Ben-David et al., 2010], one employs large unlabeled samples to estimate the relatedness of source and target domains to perform the adaptation. Even if unlabeled data are abundant, the estimation of adaptation parameters can be computationally prohibitive. This is the case, for example, when a large number of domains is involved or when one acquires new domains incrementally.

A recently proposed setting, closer to the one we consider, is the Hypothesis Transfer Learning (HTL) [Kuzborskij and Orabona, 2013, Ben-David and Uner, 2013], where the practical limitations of TL and DA are alleviated through indirect access to the *source domain* by means of a *source hypothesis*. Also, in the HTL setting there are no restrictions on how the source hypotheses can be used to boost the performance on the target task.

Albeit empirically the setting considered in this paper has already been extensively exploited in the past [Yang et al., 2007, Orabona et al., 2009, Tommasi et al., 2010, Luo et al., 2011, Kuzborskij et al., 2013], a first theoretical treatment of this setting was given by Kuzborskij and Orabona [2013], where we analyzed the linear HTL algorithm that solves a regularized least-squares problem with a single fixed, unweighted, source hypothesis. We proved a polynomial generalization bound that depends on the performance of the fixed source hypothesis on the target task.

Our contributions. We extend the formulation in [Kuzborskij and Orabona, 2013], with a general regularized Empirical Risk Minimization (ERM) problem with respect to any non-negative smooth loss function, not necessarily convex, and any strongly convex regularizer. We prove high-probability generalization bounds that exhibit *fast rate*, i.e. $\mathcal{O}(1/m)$, of convergence whenever any *weighted combination* of multiple source hypotheses performs well on the target task. In addition, we show that, if the combination is perfect, the error on the training set becomes deterministically equal to the generalization error. Furthermore, we analyze the excess risk of our formulation, and conclude that a good source hypothesis also speeds up the convergence to the performance of the best-in-the-class. As a byproduct of our study, we prove an upper bound on the Rademacher complexity of a smooth loss class that provides extra information compared to that of Lipschitz loss classes. Our analysis, that might be of independent interest, is an alternative to the analysis of Srebro et al. [2010] and it holds under much weaker assumptions.

The rest of the paper is organized as follows. In the next section we make a brief review of the previous work. Next, we formally state our formulation in Section 4 and present main results right after, in Section 5. In Section 5.1 we discuss the implications and compare them to the body of literature in learning with the fast rates and transfer learning. Next, in Section 6, we present the proofs of our main results. Section 7 concludes the paper.

2 Related Work

Kuzborskij and Orabona [2013] showed that the generalization ability of the regularized least-squares HTL algorithm improves if the supplied *source hypothesis* performs well on the target task. More specifically, we proposed a key criterion, *the risk of the source hypothesis on the target domain*, that captures the relatedness of the source and target domains. Later, Ben-David and Uner [2013] showed a similar bound, but with a different quantity capturing the relatedness between source and target. Instead of considering a general source hypothesis, they have confined their analysis to the linear hypothesis class. This allowed them to show that the target hypothesis generalizes better when it is close to the good source hypothesis. From this

perspective it is easy to interpret the source hypothesis as an initialization point in the hypothesis class. Naturally, given a starting position that is close to the best in the class, one generalizes well.

Prior to these works there were few studies trying to understand the learning with auxiliary hypotheses subject to different conditions. Li and Bilmes [2007] have analyzed a Bayesian approach to HTL. Employing a PAC-Bayes analysis they showed that given a prior on the hypothesis class, the generalization ability of logistic regression improves if the prior is informative on the target task. Mansour et al. [2008] analyzed a setting of *multiple source hypotheses* combination. There, in addition to the source hypotheses, the learner receives unlabeled samples drawn from the source distributions, that are used to weight and combine these source hypotheses. They have studied the possibility of learning in such a scenario, however, they did not address the generalization properties of any particular algorithm.

Unlike these works, we focus on the generalization ability of a large family of HTL algorithms, that generate the target predictor given a set of multiple source hypotheses. In particular, we analyze Regularized Empirical Risk Minimization with the choice of any non-negative smooth loss and any strongly convex regularizer. Thus our analysis covers a wide range of algorithms, explaining their empirical success. One category of those, prevalent in computer vision [Kienzle and Chellapilla, 2006, Yang et al., 2007, Tommasi et al., 2010, Aytar and Zisserman, 2011, Kuzborskij et al., 2013, Tommasi et al., 2014], employs the principle of biased regularization [Schölkopf et al., 2001]. For example, instead of penalizing large weights by introducing the term $\|\mathbf{w}\|^2$ into the objective function, one enforces them to be close to some “prior” model, that is $\|\mathbf{w} - \mathbf{w}^{\text{prior}}\|^2$. This principle also found its applications in other fields, such as NLP [Daumé III, 2007, Daumé III et al., 2010], and electromyography classification [Orabona et al., 2009, Tommasi et al., 2013]. Many empirical works have also investigated the use of the source hypotheses in a “black box” sense, sometimes not even posing the problem as a transfer learning [Duan et al., 2009, Li et al., 2010, Luo et al., 2011, Bergamo and Torresani, 2014], and recently in conjunction with deep neural networks [Oquab et al., 2014].

In the literature there are several other machine learning directions conceptually similar to the one we consider in this work. Arguably, the most well known one is the Domain Adaptation (DA) problem. The standard machine learning assumption is that the training and the testing sets are sampled from the same probability distribution. In such case, we expect that a hypothesis generated by the learner from that training set will lead to sensible predictions on the testing set. The difficulty arises when training and testing distributions differ, that is we have a training set sampled from the *source domain* and testing set from the *target domain*. Clearly, the hypothesis generated from the source domain can perform arbitrarily badly on the target one. A paradigm of DA, addressing this issue has received a lot of attention in recent years [Ben-David et al., 2010, Mansour et al., 2009]. Although, this framework is different from the one we study in this work, we identify similarities and compare our findings with the theory of learning from different domains in Section 5.2.

3 Definitions

In this section we introduce the definitions used in the rest of the paper.

We denote random variables by capital letters. The expected value of a random variable distributed according to a probability distribution \mathcal{D} is denoted by $\mathbb{E}_{X \sim \mathcal{D}}[X]$ and the variance is denoted by $\text{Var}_{X \sim \mathcal{D}}[X]$. The small and capital bold letters will stand respectively for the vectors and matrices, e.g. $\mathbf{x} = [x_1, \dots, x_d]^\top$ and $\mathbf{A} \in \mathbb{R}^{d_1 \times d_2}$.

Denoting by \mathcal{X} and \mathcal{Y} respectively the input and output space of the learning problem, the training set is $S = \{(\mathbf{x}_i, y_i)\}_{i=1}^m$, drawn i.i.d. from the probability distribution \mathcal{D} defined over $\mathcal{X} \times \mathcal{Y}$. Without the loss of generality we will have $\mathcal{X} = \{\mathbf{x} : \|\mathbf{x}\| \leq 1\}$ and we will focus on the problems where $\mathcal{Y} = [-C, C]$.

To measure the accuracy of a learning algorithm, we introduce a non-negative loss function $\ell(h(\mathbf{x}), y)$, which measures the cost incurred predicting $h(\mathbf{x})$ instead of y . The *risk* of a hypothesis h , with respect to

a probability distribution \mathcal{D} , and the *empirical risk* measured on the sample S are then defined as

$$R(h) := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\ell(h(\mathbf{x}), y)], \quad \text{and} \quad \hat{R}_S(h) := \frac{1}{m} \sum_{i=1}^m \ell(h(\mathbf{x}_i), y_i).$$

In the following, the risk is measured with respect to the probability distribution of the *target* domain, unless stated otherwise. We capture the smoothness of the loss function via following definition.

H -smooth loss function. We say that a non-negative loss function $\ell : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}_+$ is *H -smooth* iff,

$$\forall t, r \in \mathbb{R}, \forall y \in \mathcal{Y}, |\nabla_t \ell(t, y) - \nabla_r \ell(r, y)| \leq H|t - r|.$$

In this work we will make use of strongly convex regularizers, functions that are defined as follows.

Strongly convex function. A function Ω is σ -strongly convex w.r.t. a norm $\|\cdot\|$ iff for all \mathbf{w}, \mathbf{v} , and $\alpha \in (0, 1)$ we have

$$\Omega(\alpha \mathbf{w} + (1 - \alpha) \mathbf{v}) \leq \alpha \Omega(\mathbf{w}) + (1 - \alpha) \Omega(\mathbf{v}) - \frac{\sigma}{2} \alpha (1 - \alpha) \|\mathbf{w} - \mathbf{v}\|^2.$$

We will quantify the complexity of a hypothesis class by the means of Rademacher complexity [Bartlett and Mendelson, 2003]. In particular, the empirical Rademacher complexity of the hypothesis class \mathcal{H} measured on the sample S and its expectation are defined as

$$\hat{\mathfrak{R}}_S(\mathcal{H}) := \mathbb{E}_{\varepsilon} \left[\sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \varepsilon_i h(\mathbf{x}_i) \right] \quad \text{and} \quad \mathfrak{R}(\mathcal{H}) := \mathbb{E}_S \left[\hat{\mathfrak{R}}_S(\mathcal{H}) \right].$$

Here, ε_i is a random variable such that $\mathbb{P}(\varepsilon_i = 1) = \mathbb{P}(\varepsilon_i = -1) = \frac{1}{2}$. Similarly, as in the case of the risk, the Rademacher complexity is measured with respect to the probability distribution of the target domain, unless stated otherwise.

4 Transferring from Auxiliary Hypotheses

In the following we will capture and generalize many transfer learning formulations that employ a collection of given *source hypotheses* $\{h_i^{\text{src}} : \mathcal{X} \mapsto \mathcal{Y}\}_{i=1}^n$ within the framework of Regularized Empirical Risk Minimization (ERM). These problems typically involve a criterion for source hypothesis selection and combination with the goal to increase performance on the *target task* [Yang et al., 2007, Tommasi et al., 2014, Kuzborskij et al., 2015]. Indeed, some source hypotheses might come from tasks similar to the target task and the goal of an algorithm is to select only relevant ones. In this work we will consider source combination

$$h_{\beta}^{\text{src}}(\mathbf{x}) := \sum_{i=1}^n \beta_i h_i^{\text{src}}(\mathbf{x}),$$

and target hypothesis

$$h_{\mathbf{w}, \beta}(\mathbf{x}) := \langle \mathbf{w}, \mathbf{x} \rangle + h_{\beta}^{\text{src}}(\mathbf{x}), \tag{1}$$

with the relevance of the sources characterized by the parameter $\beta \in \mathbb{R}^n$. We will focus on the Regularized ERM formulations with the choice of any non-negative smooth loss function and any strongly-convex regularizer. This puts our problem into the class of the ones that can be solved efficiently, yet endowed with interesting properties.

Regularized ERM for Transferring from Auxiliary Hypotheses. Let $\ell : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}_+$ be an H -smooth loss function and let $\Omega : \mathcal{H} \mapsto \mathbb{R}_+$ be a σ -strongly convex function w.r.t. a norm $\|\cdot\|$. Given the target training set $S = \{(\mathbf{x}_i, y_i)\}_{i=1}^m$, $\lambda \in \mathbb{R}_+$, source hypotheses $\{h_i^{\text{src}}\}_{i=1}^n$, and parameters β obeying $\Omega(\beta) \leq \rho$, the algorithm generates the *target hypothesis* $h_{\hat{\mathbf{w}}, \beta}$, such that

$$\hat{\mathbf{w}} = \underset{\mathbf{w} \in \mathcal{H}}{\operatorname{argmin}} \left\{ \frac{1}{m} \sum_{i=1}^m \ell(\langle \mathbf{w}, \mathbf{x}_i \rangle + h_{\beta}^{\text{src}}(\mathbf{x}), y_i) + \lambda \Omega(\mathbf{w}) \right\}. \quad (2)$$

Note that (2) is minimized only w.r.t. \mathbf{w} , that is, we do not analyze any particular algorithm that searches for the optimal weights of the source hypotheses. However, we assume that $\Omega(\beta) \leq \rho$, that is we constrain β through a strongly convex function. Thus, we cover regularized algorithms generating β , which includes most of the empirical work in this field, and potential new algorithms.

In the following we will pay special attention to a quantity that captures the performance of the source hypothesis combination $h_{\beta}^{\text{src}}(\mathbf{x})$ on the target domain

$$R^{\text{src}} := R(h_{\beta}^{\text{src}}).$$

Our analysis will focus on the generalization properties of $h_{\hat{\mathbf{w}}, \beta}$. In particular, our main goal will be to understand the impact of the source hypothesis combination on the performance of the target hypothesis. In our analysis we will discuss various regimes of interest, for example considering the perfect and arbitrarily bad source hypothesis. Our discussion will touch scenarios where the auxiliary hypotheses accelerate the learning and the conditions when we can provably expect perfect generalization. Finally, we will consider the consistency of the algorithm (1) and pinpoint conditions when we achieve faster convergence to the performance of the best-in-the-class.

One special example covered by our analysis, commonly applied in transfer learning, is the *biased regularization* [Schölkopf et al., 2001]. Consider the following least-squares based algorithm.

Least-Squares with Biased Regularization. Given the target training set $S = \{(\mathbf{x}_i, y_i)\}_{i=1}^m$, source hypotheses $\{\mathbf{w}_i^{\text{src}}\}_{i=1}^n \subset \mathcal{H}$, parameters $\beta \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}_+$, the algorithm generates the target hypothesis $h(\mathbf{x}) = \langle \hat{\mathbf{w}}, \mathbf{x} \rangle$, where

$$\hat{\mathbf{w}} = \underset{\mathbf{w} \in \mathcal{H}}{\operatorname{argmin}} \left\{ \frac{1}{m} \sum_{i=1}^m (\langle \mathbf{w}, \mathbf{x}_i \rangle - y_i)^2 + \lambda \|\mathbf{w} - \mathbf{W}^{\text{src}} \beta\|_2^2 \right\}. \quad (3)$$

This problem has a simple intuitive interpretation: minimize the training error on the target training set while keeping the solution close to the linear combination of the source hypotheses. One can naturally arrive at (3) from a probabilistic perspective: The solution $\hat{\mathbf{w}}$ is a maximum a posteriori estimate when the conditional distribution is Gaussian and the prior is a $\mathbf{W}^{\text{src}} \beta$ -mean, $\frac{1}{\lambda} \mathbf{I}$ -covariance Gaussian distribution. Even though biased regularization is a simple idea, it found success in a plethora of transfer learning applications, ranging from computer vision [Kienzle and Chellapilla, 2006, Yang et al., 2007, Tommasi et al., 2010, Aytar and Zisserman, 2011, Kuzborskij et al., 2013, Tommasi et al., 2014] to NLP [Daumé III, 2007], to electromyography classification [Orabona et al., 2009, Tommasi et al., 2013].

Claim 1 *Least-Squares with Biased Regularization is a special case of the Regularized ERM in (1).*

Proof Introduce \mathbf{w}' , such that $\mathbf{w}' = \mathbf{w} - \mathbf{W}^{\text{src}} \beta$. Then we have that problem (3) is equivalent to

$$\min_{\mathbf{w} \in \mathcal{H}} \left\{ \frac{1}{m} \sum_{i=1}^m (\langle \mathbf{w}' + \mathbf{W}^{\text{src}} \beta, \mathbf{x}_i \rangle - y_i)^2 + \lambda \|\mathbf{w}'\|_2^2 \right\},$$

that in turn is a special version of (2) when $h_i^{\text{src}}(\mathbf{x}) = \langle \mathbf{w}_i^{\text{src}}, \mathbf{x} \rangle$, we use the square loss, and $\|\cdot\|_2^2$ as regularizer. \blacksquare

Albeit practically appealing, the formulation (3) is limited in the fact that the source hypotheses must be a linear predictor living in the same space of the target predictor. Instead, the formulation in (1) naturally generalizes the biased regularization formulation, allowing to treat the source hypothesis as “black box” predictors.

5 Main Results

In this section, we present the main results of this work: generalization and excess risk bounds for the Regularized ERM. In the next section we discuss in detail the implications of these results, while we defer the proofs to the subsequent sections.

The first bound demonstrates the utility of the perfect combination of source hypotheses, while the second lets us observe the dependency on the arbitrary combination. In particular, the first bound explicitates the intuition that given the perfect source hypothesis learning is not required. In other words, when $R^{\text{src}} = 0$ we have that the empirical risk becomes equal to the risk with probability one.

Theorem 2 *Let $h_{\hat{\mathbf{w}},\beta}$ be generated by Regularized ERM, given the m -sized training set S sampled i.i.d. from the target domain, source hypotheses $\{h_i^{\text{src}} : \|h_i^{\text{src}}\|_\infty \leq 1\}_{i=1}^n$, any source weights β obeying $\Omega(\beta) \leq \rho$, and $\lambda \in \mathbb{R}_+$. Assume that $\ell(h_{\hat{\mathbf{w}},\beta}(\mathbf{x}), y) \leq M$ for any (\mathbf{x}, y) and any training set. Then, denoting $\kappa = \frac{H}{\sigma}$ and assuming that $\lambda \leq \kappa$, we have with probability at least $1 - e^{-\eta}$, $\forall \eta \geq 0$*

$$R(h_{\hat{\mathbf{w}},\beta}) \leq \hat{R}_S(h_{\hat{\mathbf{w}},\beta}) + \mathcal{O} \left(\frac{R^{\text{src}} \kappa}{\sqrt{m} \lambda} + \sqrt{\frac{R^{\text{src}} \rho \kappa^2}{m \lambda}} + \frac{M \eta}{m \log \left(1 + \sqrt{\frac{M \eta}{u^{\text{src}}}} \right)} \right) \quad (4)$$

$$\leq \hat{R}_S(h_{\hat{\mathbf{w}},\beta}) + \mathcal{O} \left(\frac{\kappa}{\sqrt{m}} \left(\frac{R^{\text{src}}}{\lambda} + \sqrt{\frac{R^{\text{src}} \rho}{\lambda}} \right) + \frac{\kappa}{m} \left(\frac{\sqrt{R^{\text{src}} M \eta}}{\lambda} + \sqrt{\frac{\rho}{\lambda}} \right) \right), \quad (5)$$

where $u^{\text{src}} = R^{\text{src}} \left(m + \frac{\kappa \sqrt{m}}{\lambda} \right) + \kappa \sqrt{\frac{R^{\text{src}} m \rho}{\lambda}}$.

Now we focus on the consistency of the HTL. Specifically, we show an upper bound on the excess risk of the Regularized ERM, which depends on R^{src} , that is the risk of the combined source hypothesis h_β^{src} on the target domain. We observe that for a small R^{src} , the excess risk shrinks at a fast rate of $\mathcal{O}(1/m)$. In other words, a good prior knowledge guarantees not only good generalization, but also the fast recovery of the performance of the best hypothesis in the class.

This bound is similar in spirit to the results of localized complexities, as in works of Bartlett et al. [2005], Srebro et al. [2010], however we focus on the linear HTL scenario rather than a generic learning setting. Later, in Section 5.1, we compare our bounds to these works and show that our analysis achieves superior results.

Theorem 3 *Let $h_{\hat{\mathbf{w}},\beta}$ be generated by Regularized ERM, given the m -sized training set S sampled i.i.d. from the target domain, source hypotheses $\{h_i^{\text{src}} : \|h_i^{\text{src}}\|_\infty \leq 1\}_{i=1}^n$, any source weights β obeying $\Omega(\beta) \leq \rho$, and $\lambda \in \mathbb{R}_+$. Then, denoting $\kappa = \frac{H}{\sigma}$, assuming that $\lambda \leq \kappa \leq 1$, and setting the regularization parameter*

$$\lambda = \mathcal{O} \left(\sqrt{\frac{\kappa R^{\text{src}} + \sqrt{R^{\text{src}} \rho}}{\tau \sqrt{m}}} + \frac{\sqrt{\kappa}}{\tau} \sqrt{\frac{R^{\text{src}} + \sqrt{R^{\text{src}} \rho}}{m^{1.5}}} \right),$$

for any choice of $\tau \geq 0$, we have with high probability that

$$\begin{aligned} R(h_{\hat{\mathbf{w}},\beta}) - \min_{\Omega(\mathbf{w}) \leq \tau} R(h_{\mathbf{w},\beta}) \\ = \mathcal{O} \left(\frac{\sqrt{R^{\text{src}}} + \sqrt[4]{R^{\text{src}}\rho}}{\sqrt[4]{m}} \sqrt{\kappa\tau} + \frac{\sqrt[4]{R^{\text{src}}} + \sqrt[8]{R^{\text{src}}\rho}}{\sqrt[4]{m^{1.5}}} \sqrt[4]{\kappa\tau^2} + \sqrt{\frac{R^{\text{src}}}{m}} + \frac{1}{m} \right). \end{aligned}$$

5.1 Implications

We start by discussing the effect on the generalization ability of the source hypothesis combination. Intuitively, a good source hypothesis combination should facilitate transfer learning, while a reasonable algorithm must not fail if we provide it with the bad one. That said, a natural question to ask here is, what makes a good or bad source hypothesis? As in previous works in transfer learning and domain adaptation, we capture this notion via a quantity that has two-fold interpretation: (1) the performance of the source hypothesis combination on the target domain; (2) relatedness of source and target domains. In the theorems presented in the previous sections we denoted it by R^{src} , that is the risk of the source hypothesis combination on the target domain. In this section we will consider various regimes of interest with respect to R^{src} .

When the source is a bad fit. First consider the case when the source hypothesis combination h_{β}^{src} is useless for the purpose of transfer learning, for example, $h_{\beta}^{\text{src}}(\mathbf{x}) = 0$ for all \mathbf{x} . This corresponds to learning with no auxiliary information. Then we can assume that $R^{\text{src}} \leq M$, and from Theorem 2 we obtain $R(h_{\hat{\mathbf{w}}}) - \hat{R}_S(h_{\hat{\mathbf{w}}}) \leq \mathcal{O}(1/(\sqrt{m}\lambda))$. This rate matches the one in the analysis of regularized least-squares [Vito et al., 2005, Bousquet and Elisseeff, 2002], that is a special case of the smooth loss function that the Regularized ERM employs. On the other hand, Srebro et al. [2010] showed a better worst-case rate $\mathcal{O}(1/\sqrt{m\lambda})$. However, their framework builds upon a worst case Rademacher complexity which does not involve the expectation over the sample and does not lead to the dependency on R^{src} we have obtained in Theorem 2. We will discuss this problem in details later.

When the source is a good fit. Here we would like to consider the behavior of the algorithm in the finite-sample and asymptotic scenarios. We first look at the regime of small m , in particular $m = \mathcal{O}(1/R^{\text{src}})$. In this case, the fast rate term will dominate the bound, and we obtain the convergence rate of $\mathcal{O}(\sqrt{\rho}/(m\sqrt{\lambda}))$. In other words, we can expect a faster convergence when m is small, where “small” depends on R^{src} , the quality of combined source hypotheses. Now consider the asymptotic behavior of the algorithm, particularly when m goes to infinity. In such case, the algorithm exhibits a rate of $\mathcal{O}\left(R^{\text{src}}/\sqrt{m\lambda} + \sqrt{(R^{\text{src}}\rho)/m\lambda}\right)$, so R^{src} controls the constant factor of the rate. Hence, the quantity R^{src} governs the transient regime for small m and the asymptotic behavior of the algorithm, predicting a faster convergence in both regimes when it is small.

When source is a perfect fit. It is conceivable that the source hypothesis exploited is the perfect one, that is $R^{\text{src}} = 0$. In other words, the source hypothesis combination is a perfect predictor for the target domain. Theorem 2 implies that $R(h_{\hat{\mathbf{w}},\beta}) = \hat{R}_S(h_{\hat{\mathbf{w}},\beta})$ with probability one. We note that for many practically used smooth losses, such as square loss, this setting is only realistic if source and target domains match and the problem is noise-free. However, we can observe $R^{\text{src}} = 0$, for example, when the squared hinge loss, $\ell(z, y) = \max\{0, 1 - zy\}^2$, is used and all target domain examples are classified correctly by the source hypothesis combination, case that is not unthinkable for related domains.

Fast rates. There is a number of works in the literature investigating a rate of convergence faster than $1/\sqrt{m}$ subject to different conditions. In particular, the localized Rademacher complexity bounds of Bartlett et al. [2005] and Bousquet [2002] can be used to obtain results similar to the second inequality of Theorem 2. Indeed, Theorem 7 shows a bound which is very similar to the localized ones, albeit with two differences. The r.h.s. of the first inequality in Theorem 7 vanishes when the loss class has zero variance. Though intuitively trivial, this allows to prove a considerable result in the theory of transfer learning as it quantifies the intuition that no learning is necessary if the source has perfect performance on the target task. Second, by applying the standard localized Rademacher complexity bounds of Bousquet [2002], and

assuming the use of the Lipschitz loss function, we do not achieve a fast rate of convergence, as can be seen from Theorem 13, shown in the Appendix. We suspect that assuming the smoothness of the loss function is crucial to prove fast rates in our formulation.

Fast rates for ERM with the smooth loss have been thoroughly analyzed by Srebro et al. [2010]. Yet, the analysis of our HTL algorithm within their framework would yield a bound that is inferior to ours in two respects. The first concerns the scenario when the combined source hypothesis is perfect, that is $R^{\text{src}} = 0$. The generalization bound of Srebro et al. [2010] does not offer a way to show that the empirical risk converges to the risk with probability one – instead one can only hope to get a fast rate of convergence. The second problem is in the fact that such bound would depend on the empirical performance of combined source hypothesis. As we have noted before, the quantity R^{src} is essential because it captures the degree of relatedness between two domains. In their bounds, one cannot obtain this relationship through the Rademacher complexity term as we did in our analysis. The reason for this is the stronger notion of Rademacher complexity that is employed by that framework, involving a supremum over the sample instead of an expectation. The expectation over the sample of the target distribution is crucial here, because it allows us to quantify how well the source domain is aligned with the target domain, through the source hypothesis acting as a link. However, one can attempt to obtain the bound on the empirical risk in terms of R^{src} . We prove such a bound in the Appendix, Theorem 11, and conclude that if one has a good source hypothesis or even a perfect one, the rate is $\mathcal{O}(1/\sqrt[4]{m^3})$, which is worse than ours.

5.2 Comparison to Theories of Domain Adaptation and Transfer Learning

The setting in DA is different from the one we study, however, we will briefly discuss the theoretical relationship between the two. Typically in DA, one trains a hypothesis from an altered source training set, striving to achieve good performance on the target domain. The key question here is how to alter, or to *adapt*, the source training set. To answer this question, DA literature introduces the notion of domain relatedness, which quantifies the dissimilarities between the marginal distributions of corresponding domains. Practically, in some cases the domain relatedness can be estimated through a large set of unlabeled samples drawn from both source and target domains. Theories of DA [Ben-David et al., 2010, Mansour et al., 2009, Ben-David and Uner, 2012, Mansour et al., 2008, Cortes and Mohri, 2014] have proposed a number of such domain relatedness criteria. Perhaps the most well known are the $d_{\mathcal{H}\Delta\mathcal{H}}$ -divergence [Ben-David et al., 2010] and its more general counterpart, the Discrepancy Distance [Mansour et al., 2009]. Typically, this divergence is explicated in the generalization bound along with other terms controlling the generalization on the target domain. Let $R_{\mathcal{D}^{\text{trg}}}(h)$ and $R_{\mathcal{D}^{\text{src}}}(h)$ denote the risks of the hypothesis h , measured w.r.t. the target and source distributions. Then a well-known result of Ben-David et al. [2010] suggests that for all $h \in \mathcal{H}$

$$R_{\mathcal{D}^{\text{trg}}}(h) \leq R_{\mathcal{D}^{\text{src}}}(h) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^{\text{src}}, \mathcal{D}^{\text{trg}}) + \varepsilon_{\mathcal{H}}^*, \quad (6)$$

where $\varepsilon_{\mathcal{H}}^* = \min_{h \in \mathcal{H}} \{R_{\mathcal{D}^{\text{trg}}}(h) + R_{\mathcal{D}^{\text{src}}}(h)\}$. This result implies that adaptation is possible given that $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^{\text{src}}, \mathcal{D}^{\text{trg}})$ and ε^* are small. One can try to reduce those by controlling the complexity of the class \mathcal{H} and by minimizing the divergence $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^{\text{src}}, \mathcal{D}^{\text{trg}})$. In practice, the latter can be manipulated through an empirical counterpart on the basis of unlabeled samples. Increasing the complexity of \mathcal{H} indeed reduces ε^* , but inflates $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^{\text{src}}, \mathcal{D}^{\text{trg}})$. On the other hand, by minimizing $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^{\text{src}}, \mathcal{D}^{\text{trg}})$ alone puts us under the risk of increasing ε^* , since the empirical divergence is reduced without taking the labelling into account.

Clearly, this bound cannot be directly compared to our result, Theorem 2. However, we note the term R^{src} appearing in our results, which plays a role very similar to $d_{\mathcal{H}\Delta\mathcal{H}}$ in (6). In fact, by defining $\mathcal{H} = \{\mathbf{x} \mapsto \langle \boldsymbol{\beta}, \mathbf{h}^{\text{src}}(\mathbf{x}) \rangle : \Omega(\boldsymbol{\beta}) \leq \tau\}$, where $\mathbf{h}^{\text{src}}(\mathbf{x}) = [h_1^{\text{src}}(\mathbf{x}), \dots, h_n^{\text{src}}(\mathbf{x})]^\top$, and fixing $h = h_{\boldsymbol{\beta}}^{\text{src}} \in \mathcal{H}$ in (6), we can write

$$R^{\text{src}} = R_{\mathcal{D}^{\text{trg}}}(h_{\boldsymbol{\beta}}^{\text{src}}) \leq R_{\mathcal{D}^{\text{src}}}(h_{\boldsymbol{\beta}}^{\text{src}}) + d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^{\text{src}}, \mathcal{D}^{\text{trg}}) + \varepsilon_{\mathcal{H}}^*.$$

Plugging this into the generalization bound (5) and assuming that $\lambda \leq 1$ and $\rho \leq 1/\lambda$ we have for the target hypothesis h that

$$R_{\mathcal{D}^{\text{trg}}}(h) \leq \hat{R}_S(h) + \mathcal{O} \left(\frac{R_{\mathcal{D}^{\text{src}}}(h_{\beta}^{\text{src}}) + d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^{\text{src}}, \mathcal{D}^{\text{trg}}) + \epsilon_{\mathcal{H}}^*}{\sqrt{m\lambda}} + \frac{1}{m\lambda} \right). \quad (7)$$

Albeit this inequality shows the generalization ability of the transfer learning algorithm, comparing to (6), we observe that DA and our result agree on the fact that the divergence between the domains has to be small to generalize well. In fact, in the formulation we consider, the divergence is controlled in two ways: implicitly, by the choice of h^{src} and through the complexity of class \mathcal{H} , that is by choosing τ . Second, in DA we expect that a hypothesis performs well on the target only if it performs well on the source. In our results, this requirement is relaxed. As a side note, we observe that (7) captures an intuitive notion that a good source hypothesis has to perform well on its own domain. Finally, in the theory of DA $\epsilon_{\mathcal{H}}^*$ is assumed to be small. Indeed, if $\epsilon_{\mathcal{H}}^*$ is large, there is no hypothesis that is able to perform well on both domains simultaneously, and therefore adaptation is hopeless. In our case, the algorithm can still generalize even with large $\epsilon_{\mathcal{H}}^*$, however this is due to the supervised nature of the framework.

We now turn our attention to the previous theoretical works studying HTL-related settings. Few papers have addressed the theory of transfer learning, where the only information passed from the source domain is the classifier or regressor. Mansour et al. [2008] have addressed the problem of multiple source hypotheses combination, however, in a different setting. Specifically, in addition to the source hypotheses, the learner receives the unlabeled samples drawn from the source distributions, that are used to weight and combine these source hypotheses. The authors have presented a general theory of such a scenario and did not study the generalization properties of any particular algorithm. The first analysis of the generalization ability of HTL in the similar context we consider here was done by Kuzborskij and Orabona [2013]. The work focused on the L_2 -regularized least squares and the generalization bound involving the leave-one-out risk instead of the empirical one. The following result, obtained through an algorithmic stability argument [Bousquet and Elisseeff, 2002], holds with probability at least $1 - \delta$

$$R(h) \leq \hat{R}_S^{\text{loo}}(h) + \mathcal{O} \left(\frac{\sqrt[4]{R^{\text{src}}}}{\sqrt{m\delta\lambda^{0.75}}} \right), \quad (8)$$

where R^{src} is the risk of a single fixed source hypothesis and h is the solution of a Regularized Least Square problem. We first observe that the shape of the bound is similar to the one obtained in this work, although with the number of differences. First, contrary to our presented bounds, their bound assumes the use of a fixed source hypothesis, that is not even weighted by any coefficient. In practice, this is a very strong assumption, as one can receive an arbitrarily bad source and have no way to exclude it. Second, the bound (8) seems to have a vanishing behavior whenever the risk of the source R^{src} is equal to zero. This comes at the cost of the use of a weaker concentration inequality. In Theorem 2 we manage to obtain the same behavior with high probability. Finally, we get a better dependency on R^{src} .

5.3 Combining Source Hypotheses in Practice

So far we have assumed that the problem (4) is supplied with a pre-made combination of source hypotheses, that is, we did not study a particular algorithm for tuning the β weights. However, by analyzing our generalization bound (2), it is easy to come up with algorithms that could be used for this purpose. In particular, by minimizing the bound w.r.t. β , and assuming that the empirical risk $\hat{R}_S(h_{\beta}^{\text{src}})$ converges uniformly to R^{src} , we have with high probability that

$$\min_{\Omega(\beta) \leq \rho} R(h_{\hat{\mathbf{w}}, \beta}) \leq \min_{\Omega(\beta) \leq \rho} \left\{ \hat{R}_S(h_{\hat{\mathbf{w}}, \beta}) + \mathcal{O} \left(\frac{\kappa \hat{R}_S(h_{\beta}^{\text{src}})}{\sqrt{m\lambda}} + \sqrt{\frac{\kappa^2 \rho \hat{R}_S(h_{\beta}^{\text{src}})}{m\lambda}} \right) \right\}.$$

Thus, at least theoretically, given a fixed solution $\hat{\mathbf{w}}$, it is enough to jointly minimize the error of the target hypothesis $h_{\hat{\mathbf{w}},\beta}$ and the error of the source combination on the target training set. This is particularly efficient when the square loss is used, since $\hat{\mathbf{w}}$ can be expressed in terms of an inverse of a covariance matrix that has to be inverted only once [Orabona et al., 2009, Tommasi et al., 2014, Kuzborskij et al., 2015].

Many HTL-like algorithms can be captured through the above by choosing among different loss functions and regularizers Ω . The simplest case is just a concatenation of the source hypotheses predictions with the original feature vector. However, by choosing different regularizers and their parameters, we can treat the source hypotheses in a different way from the original features. For example, one might enforce sparsity over the source hypotheses, while using the usual L2 regularizer on the target solution $\hat{\mathbf{w}}$.

6 Technical Results and Proofs

In this section we present general technical results that are used to prove our theorems.

First, we present the Rademacher complexity generalization bound in Theorem 7, which slightly differs from the usual ones. The difference comes in the assumption that the variance of the loss is uniformly bounded over the hypothesis class. This will allow us to state a generalization bound that obeys the fast empirical risk convergence rate subject to the small class complexity. Second, we will also show a generalization bound with the confidence term that vanishes if the complexity of the class is exactly zero.

Next, we focus on the Rademacher complexity of the smooth loss function class. We prove a bound on the empirical Rademacher complexity of a hypothesis class, Lemma 8, that depends on the point-wise bounds on the loss function. This novel bound might be of independent interest.

Finally, we employ this result to analyze the effect of the source hypotheses on the complexity of the target hypothesis class in Theorem 10.

6.1 Fast Rate Generalization Bound

The proof of fast-rate and vanishing-confidence-term bounds, Theorem 7, stems from the functional generalization of Bennett’s inequality which is due to Bousquet [2002, Theorem 2.11] and that we report here for completeness.

Theorem 4 (Bousquet [2002]) *Let X_1, X_2, \dots, X_m be identically distributed random variables according to \mathcal{D} . For all \mathcal{D} -measurable, square-integrable $g \in \mathcal{G}$, with $\mathbb{E}_X[g(X)] = 0$, and $\sup_{g \in \mathcal{G}} \text{ess sup } g \leq 1$, we denote*

$$Z = \sup_{g \in \mathcal{G}} \sum_{i=1}^m g(X_i). \quad (9)$$

Let σ be a positive real number such that $\sup_{g \in \mathcal{G}} \text{Var}_{X \sim \mathcal{D}}[g(X)] \leq \sigma^2$ almost surely. Then for all $t \geq 0$, we have that

$$\mathbb{P}(Z \geq \mathbb{E}[Z] + t) \leq \exp\left(-vu \left(\frac{t}{v}\right)\right), \quad (10)$$

where

$$v = m\sigma^2 + 2\mathbb{E}[Z],$$

$$u(y) = (1 + y) \log(1 + y) - y.$$

The following technical lemma will be used to invert the right hand side of (10).

Lemma 5 *Let $a, b > 0$ such that $b = (1 + a) \log(1 + a) - a$. Then $a \leq \frac{3b}{2 \log(\sqrt{b} + 1)}$.*

Proof It is easy to verify that the inverse function $f^{-1}(b)$ of $f(a) := (1+a)\log(1+a) - a$ is

$$f^{-1}(b) = \exp \left[W \left(\frac{b-1}{e} \right) + 1 \right] - 1,$$

where the function $W : \mathbb{R}_+ \rightarrow \mathbb{R}$ is the Lambert function that satisfies

$$x = W(x) \exp(W(x)).$$

Hence, to obtain an upper bound to a , we need an upper bound to the Lambert function. We use Theorem 2.3 in [Hoorfar and Hassani, 2008], that says that

$$W(x) \leq \log \frac{x+C}{1+\log(C)}, \quad \forall x > -\frac{1}{e}, C > \frac{1}{e}.$$

Setting $C = \frac{\sqrt{b}+1}{e}$, we obtain

$$a = f^{-1}(b) \leq e \frac{\frac{b-1}{e} + \frac{\sqrt{b}+1}{e}}{1 + \log(\frac{\sqrt{b}+1}{e})} - 1 = \frac{b + \sqrt{b}}{\log(\sqrt{b}+1)} - 1 \leq \frac{3b}{2 \log(\sqrt{b}+1)},$$

where in the last inequality we used the fact that $x + \sqrt{x} - \log(\sqrt{x}+1) \leq \frac{3}{2}x, \forall x \geq 0$, as it can be easily verified comparing the derivatives of both terms. \blacksquare

The following lemma is a standard tool [Mohri et al., 2012, (3.8)-(3.13)], [Bartlett and Mendelson, 2003].

Lemma 6 (Symmetrization) For any $f \in \mathcal{F}$, given random variables $S = \{X_i\}_{i=1}^m$, we have

$$\begin{aligned} \mathbb{E}_S \sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_X[f(X)] - \frac{1}{m} \sum_{i=1}^m f(X_i) \right\} &\leq 2\mathfrak{R}(\mathcal{F}), \\ \mathbb{E}_S \sup_{f \in \mathcal{F}} \left\{ \frac{1}{m} \sum_{i=1}^m f(X_i) - \mathbb{E}_X[f(X)] \right\} &\leq 2\mathfrak{R}(\mathcal{F}). \end{aligned}$$

Now we are ready to present the proof of Theorem 7.

Theorem 7 Consider the non-negative loss function $\ell : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}_+$, such that $0 \leq \ell(h(\mathbf{x}), y) \leq M$ for any $h \in \mathcal{H}$ and any $(\mathbf{x}, y) \in \mathcal{X} \times \mathcal{Y}$. In addition, let the training set S of size m be sampled i.i.d. from the probability distribution over $\mathcal{X} \times \mathcal{Y}$. Also for any $r \geq 0$, define the loss class with respect to the hypothesis class \mathcal{H} as,

$$\mathcal{L} := \{(\mathbf{x}, y) \mapsto \ell(h(\mathbf{x}), y) : h \in \mathcal{H} \wedge R(h) \leq r\}.$$

Then we have for all $h \in \mathcal{H}$, and any training set S of size m , with probability at least $1 - e^{-\eta}$, $\forall \eta \geq 0$

$$R(h) - \hat{R}_S(h) \leq 2\mathfrak{R}(\mathcal{L}) + \frac{3M\eta}{m \log \left(1 + \sqrt{\frac{2M\eta}{vm}} \right)} \leq 2\mathfrak{R}(\mathcal{L}) + 3\sqrt{\frac{vM\eta}{2m}} + \frac{3M\eta}{2m},$$

where $v = 4\mathfrak{R}(\mathcal{L}) + r$.

Proof To prove the statement, we will consider the uniform deviations of the empirical risk. Namely, we will show an upper bound on the random variable $\sup_{h \in \mathcal{H}} \{R(h) - \hat{R}_S(h)\}$. For this purpose, we will use the functional generalization of Bennett's inequality given by Theorem 4. Consider the random variable

$$Z := \frac{m}{2M} \sup_{h \in \mathcal{H}} \{R(h) - \hat{R}_S(h)\}.$$

Using Theorem 4, we have

$$\begin{aligned} \mathbb{P} \left(\frac{m}{2M} \sup_{h \in \mathcal{H}} \{R(h) - \hat{R}_S(h)\} \geq \frac{m}{2M} \mathbb{E} \left[\sup_{h \in \mathcal{H}} \{R(h) - \hat{R}_S(h)\} \right] + t \right) \\ \leq \exp \left(-vu \left(\frac{t}{v} \right) \right), \end{aligned} \quad (11)$$

where,

$$\begin{aligned} v &= m\sigma^2 + \frac{m}{M} \mathbb{E} \left[\sup_{h \in \mathcal{H}} \{R(h) - \hat{R}_S(h)\} \right], \\ \sigma^2 &\geq \sup_{h \in \mathcal{H}} \text{Var}_{(\mathbf{x}, y)} \left[\frac{1}{2M} \left(\ell(h(\mathbf{x}), y) - \mathbb{E}_{(\mathbf{x}', y')} [\ell(h(\mathbf{x}'), y')] \right) \right]. \end{aligned} \quad (12)$$

We now need two things: invert the r.h.s. of (11), treating it as a function of t , and provide an upper-bound on v . For the first part, recall that $u(y) = (1+y) \log(1+y) - y$. To give an upper-bound of t , we apply Lemma 5 with $a = \frac{t}{v}$, and $b = \frac{1}{v}\eta$. This leads to the inequalities

$$\frac{t}{v} \leq \frac{3\eta}{2v \log(1 + \sqrt{\frac{\eta}{v}})} \leq \frac{3\eta}{4v} + \frac{3}{2} \sqrt{\frac{\eta}{v}}.$$

Using this fact, we have with probability at least $1 - e^{-\eta}$ with any $\eta \geq 0$

$$\frac{m}{2M} \sup_{h \in \mathcal{H}} \{R(h) - \hat{R}_S(h)\} \leq \frac{m}{2M} \mathbb{E} \left[\sup_{h \in \mathcal{H}} \{R(h) - \hat{R}_S(h)\} \right] + \frac{3\eta}{2 \log(1 + \sqrt{\frac{\eta}{v}})} \quad (13)$$

$$\leq \frac{m}{2M} \mathbb{E} \left[\sup_{h \in \mathcal{H}} \{R(h) - \hat{R}_S(h)\} \right] + \frac{3}{4}\eta + \frac{3}{2}\sqrt{v\eta}. \quad (14)$$

Next we prove the bound on v . We first show that the variance of centered loss function, σ^2 , is uniformly bounded by the Rademacher complexity. From the definition of variance we have

$$\begin{aligned} \sup_{h \in \mathcal{H}} \mathbb{E}_{(\mathbf{x}, y)} \left[\frac{1}{4M^2} \left(\ell(h(\mathbf{x}), y) - \mathbb{E}_{(\mathbf{x}', y')} [\ell(h(\mathbf{x}'), y')] \right)^2 \right] &\leq \sup_{h \in \mathcal{H}} \frac{1}{4M^2} \mathbb{E}_{(\mathbf{x}, y)} [\ell(h(\mathbf{x}), y)^2] \\ &\leq \sup_{h \in \mathcal{H}} \frac{1}{2M} \mathbb{E}_{(\mathbf{x}, y)} [|\ell(h(\mathbf{x}), y)|] = \sigma^2 = \sup_{h \in \mathcal{H}} \frac{1}{2M} R(h) = \frac{r}{2M}. \end{aligned} \quad (15)$$

Last inequality is due to the fact that $\ell(h(\mathbf{x}), y) \leq M$. Now we upper-bound the second term of v by applying Lemma 6,

$$\begin{aligned} \frac{1}{2mM} \mathbb{E}_S \left[\sup_{h \in \mathcal{H}} \sum_{i=1}^m \left(\ell(h(\mathbf{x}_i), y_i) - \mathbb{E}_{(\mathbf{x}', y')} [\ell(h(\mathbf{x}'), y')] \right) \right] \\ = \frac{1}{2M} \mathbb{E}_S \left[\sup_{h \in \mathcal{H}} \left\{ \left(\frac{1}{m} \sum_{i=1}^m \ell(h(\mathbf{x}_i), y_i) \right) - \mathbb{E}_{(\mathbf{x}', y')} [\ell(h(\mathbf{x}'), y')] \right\} \right] \leq \frac{1}{M} \mathfrak{R}(\mathcal{L}). \end{aligned}$$

We conclude the proof by upper-bounding the expectation terms in (13) and (14) using Lemma 6, and plugging the upper bound on v ,

$$v \leq \frac{2m}{M} \mathfrak{R}(\mathcal{L}) + m\sigma^2 \leq \frac{2m\mathfrak{R}(\mathcal{L})}{M} + \frac{mr}{2M}.$$

■

6.2 Rademacher Complexity of Smooth Loss Class

In this section we study the Rademacher complexity of the hypothesis class populated by functions of the form (1), where the parameters \mathbf{w} and β are chosen by an algorithm with a strongly convex regularizer. For this purpose we employ the results of Kakade et al. [2008, 2012], who studied strongly convex regularizers in a more general setting. Furthermore, we will focus on the use of smooth loss functions as done by Srebro et al. [2010].

The proof of the main result of this section, Theorem 10, depends essentially on the following lemma, that bounds the empirical Rademacher complexity of a H -smooth loss class.

Lemma 8 *Let $\ell : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}_+$ be the H -smooth loss function. Then for some function class \mathcal{F} , let the loss class be*

$$\mathcal{L} = \{(\mathbf{x}, y) \mapsto \ell(f(\mathbf{x}), y) : f \in \mathcal{F}\}.$$

Then having the sample S of size m and the set

$$\{\tau_i : \tau_i \geq \ell(f(\mathbf{x}_i), y_i), \forall (\mathbf{x}_i, y_i) \in S \wedge \forall f \in \mathcal{F}\},$$

we have that

$$\hat{\mathfrak{R}}_S(\mathcal{L}) \leq \mathbb{E}_{\boldsymbol{\varepsilon}} \left[\sup_{f \in \mathcal{F}} \left\{ \frac{2\sqrt{3H}}{m} \sum_{i=1}^m \varepsilon_i \sqrt{\tau_i} f(\mathbf{x}_i) \right\} \right],$$

where ε_i is r.v. such that $\mathbb{P}(\varepsilon_i = 1) = \mathbb{P}(\varepsilon_i = -1) = \frac{1}{2}$.

Proof This proof follows a line of reasoning similar to the proof of Talagrand's lemma for Lipschitz functions, see for instance Mohri et al. [2012, p. 79]. We will also use Lemma B.1 by Srebro et al. [2010] (arXiv extended version), stating that for any H -smooth non-negative function $\phi : \mathbb{R} \mapsto \mathbb{R}_+$ and any $x, z \in \mathbb{R}$,

$$|\phi(x) - \phi(z)| \leq \sqrt{6H(\phi(x) + \phi(z))}|x - z|. \quad (16)$$

Fix the sample S , then, by definition,

$$\begin{aligned} \hat{\mathfrak{R}}_S(\mathcal{L}) &= \frac{1}{m} \mathbb{E}_{\boldsymbol{\varepsilon}} \left[\sup_{f \in \mathcal{F}} \left\{ \sum_{i=1}^m \varepsilon_i \ell(f(\mathbf{x}_i), y_i) \right\} \right] \\ &= \mathbb{E}_{\varepsilon_1, \dots, \varepsilon_{m-1}} \left[\mathbb{E}_{\varepsilon_m} \left[\sup_{f \in \mathcal{F}} \{u_{m-1}(f) + \varepsilon_m \ell(f(\mathbf{x}_m), y_m)\} \right] \right], \end{aligned}$$

where $u_{m-1}(f) = \sum_{i=1}^{m-1} \varepsilon_i \ell(f(\mathbf{x}_i), y_i)$. By definition of supremum, for any $\delta > 0$, there exist $f_1, f_2 \in \mathcal{F}$ such that

$$\begin{aligned} u_{m-1}(f_1) + \ell(f_1(\mathbf{x}_m), y_m) &\geq (1 - \delta) \left(\sup_{f \in \mathcal{F}} \{u_{m-1}(f) + \ell(f(\mathbf{x}), y)\} \right) \\ \text{and } u_{m-1}(f_2) - \ell(f_2(\mathbf{x}_m), y_m) &\geq (1 - \delta) \left(\sup_{f \in \mathcal{F}} \{u_{m-1}(f) - \ell(f(\mathbf{x}), y)\} \right). \end{aligned}$$

Thus for any $\delta > 0$, by definition of $\mathbb{E}_{\varepsilon_m}$,

$$\begin{aligned}
& (1 - \delta) \mathbb{E}_{\varepsilon_m} \left[\sup_{f \in \mathcal{F}} \{u_{m-1}(f) + \varepsilon_m \ell(f(\mathbf{x}_m), y_m)\} \right] \\
&= \frac{1 - \delta}{2} \left(\sup_{f \in \mathcal{F}} \{u_{m-1}(f) + \ell(f(\mathbf{x}_m), y_m)\} + \sup_{f \in \mathcal{F}} \{u_{m-1}(f) - \ell(f(\mathbf{x}_m), y_m)\} \right) \\
&\leq \frac{1}{2} \left(u_{m-1}(f_1) + \ell(f_1(\mathbf{x}_m), y_m) + u_{m-1}(f_2) - \ell(f_2(\mathbf{x}_m), y_m) \right) \\
&\leq \frac{1}{2} \left(u_{m-1}(f_1) + u_{m-1}(f_2) \right. \\
&\quad \left. + s_m \sqrt{6H\ell(f_1(\mathbf{x}_m), y_m) + \ell(f_2(\mathbf{x}_m), y_m)} (f_1(\mathbf{x}_m) - f_2(\mathbf{x}_m)) \right) \\
&\leq \frac{1}{2} \left(u_{m-1}(f_1) + u_{m-1}(f_2) + s_m \sqrt{12H\tau_m} (f_1(\mathbf{x}_m) - f_2(\mathbf{x}_m)) \right) \\
&\leq \frac{1}{2} \sup_{f \in \mathcal{F}} \left\{ u_{m-1}(f) + s_m \sqrt{12H\tau_m} f(\mathbf{x}_m) \right\} \\
&\quad + \frac{1}{2} \sup_{f \in \mathcal{F}} \left\{ u_{m-1}(f) - s_m \sqrt{12H\tau_m} f(\mathbf{x}_m) \right\} \\
&= \mathbb{E}_{\varepsilon_m} \left[\sup_{f \in \mathcal{F}} \left\{ u_{m-1}(f) + \varepsilon_m \sqrt{12H\tau_m} f(\mathbf{x}_m) \right\} \right].
\end{aligned}$$

To obtain the second inequality, we applied (16), where $s_m = \text{SGN}(f_1(\mathbf{x}_m) - f_2(\mathbf{x}_m))$. Since the inequality holds for all $\delta > 0$, we have

$$\mathbb{E}_{\varepsilon_m} \left[\sup_{f \in \mathcal{F}} \{u_{m-1}(f) + \varepsilon_m \ell(f(\mathbf{x}_m), y_m)\} \right] \leq \mathbb{E}_{\varepsilon_m} \left[\sup_{f \in \mathcal{F}} \left\{ u_{m-1}(f) + \varepsilon_m \sqrt{12H\tau_m} f(\mathbf{x}_m) \right\} \right].$$

Proceeding in the same way for all the other ε_i , with $i \neq m$, proves the lemma. \blacksquare

To prove Theorem 10 we will also use the following lemma in Kakade et al. [2012, Corollary 4].

Lemma 9 (Kakade et al. [2012]) *If Ω is σ strongly convex w.r.t. $\|\cdot\|$ and $\Omega^*(\mathbf{0}) = 0$, then, denoting the partial sum $\sum_{j \leq i} \mathbf{v}_j$ by $\mathbf{v}_{1:i}$, we have for any sequence $\mathbf{v}_1, \dots, \mathbf{v}_m$ and for any \mathbf{u} ,*

$$\sum_{i=1}^m \langle \mathbf{v}_i, \mathbf{u} \rangle - \Omega(\mathbf{u}) \leq \Omega^*(\mathbf{v}_{1:m}) \leq \sum_{i=1}^m \langle \nabla \Omega^*(\mathbf{v}_{1:i-1}), \mathbf{v}_i \rangle + \frac{1}{2\sigma} \sum_{i=1}^m \|\mathbf{v}_i\|_*^2.$$

Now we are ready to give the proofs of the Rademacher complexity results.

Theorem 10 *Let Ω be a non-negative σ -strongly convex function w.r.t. a norm $\|\cdot\|$, and let $\mathbf{0}$ be its minimizer. Let risk and empirical risk be defined w.r.t. an H -smooth loss function $\ell : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}_+$. Finally, given the set of functions $\{f_i : \mathcal{X} \mapsto \mathcal{Y}\}_{i=1}^n$ with $\mathbf{f}(\mathbf{x}) := [f_1(\mathbf{x}), \dots, f_n(\mathbf{x})]^\top$, a combination $f_\beta(\mathbf{x}) = \langle \beta, \mathbf{f}(\mathbf{x}) \rangle$, a scalar $\alpha > 0$, and any sample S drawn i.i.d. from distribution over $\mathcal{X} \times \mathcal{Y}$, define classes*

$$\mathcal{W} = \left\{ \mathbf{w} : \Omega(\mathbf{w}) \leq \alpha \hat{R}_S(f_\beta) \right\}, \quad \mathcal{V} = \left\{ \beta : \Omega(\beta) \leq \rho \right\},$$

and the loss class

$$\mathcal{L} = \left\{ (\mathbf{x}, y) \mapsto \ell(\langle \mathbf{w}, \mathbf{x} \rangle + f_\beta(\mathbf{x}), y) : \mathbf{w} \in \mathcal{W} \wedge \beta \in \mathcal{V} \right\}.$$

Then for the loss class \mathcal{L} , setting constants $\sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_* \leq B$ and $\sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{f}(\mathbf{x})\|_* \leq C$, we have that

$$\mathfrak{R}(\mathcal{L}) \leq 4\sqrt{3H}(B+C) \left(1 + \sqrt{\frac{2HB^2\alpha}{\sigma}}\right) \frac{R(f_\beta)\sqrt{\alpha} + \sqrt{R(f_\beta)\rho}}{\sqrt{m\sigma}}.$$

Proof The core of the proof consists in an application of Lemma 8. In particular, Lemma 8 allows us to introduce additional information about the loss class by providing bounds on the loss at each example. We will bound the loss at each example using the definition of smoothness, extracting the empirical risk of hypothesis $\hat{R}_S(f_\beta)$. The last step is to give an upper-bound on the empirical Rademacher complexity of a class regularized by a strongly convex function. We follow the proof of Kakade et al. [2012, Theorem 7] to accomplish this task. First define the classes

$$\mathcal{H}_\mathcal{W} := \{\mathbf{x} \mapsto \langle \mathbf{w}, \mathbf{x} \rangle : \mathbf{w} \in \mathcal{W}\}, \quad \mathcal{H}_\mathcal{V} := \{f_\beta : \beta \in \mathcal{V}\},$$

and also define altered samples $S' := \{\sqrt{\tau_i}\mathbf{x}_i\}_{i=1}^m$ and $S'' := \{\sqrt{\tau_i}\mathbf{f}(\mathbf{x}_i)\}_{i=1}^m$, where τ_i is a quantity independent from \mathcal{W} and \mathcal{V} . Then by applying Lemma 8, we have that,

$$\begin{aligned} \hat{\mathfrak{R}}_S(\mathcal{L}) &\leq \mathbb{E}_\varepsilon \left[\sup_{\substack{\mathbf{w} \in \mathcal{W} \\ \beta \in \mathcal{V}}} \left\{ \frac{2\sqrt{3H}}{m} \sum_{i=1}^m \varepsilon_i \sqrt{\tau_i} \langle \mathbf{w}, \mathbf{x}_i \rangle + \frac{2\sqrt{3H}}{m} \sum_{i=1}^m \varepsilon_i \sqrt{\tau_i} \langle \beta, \mathbf{f}(\mathbf{x}_i) \rangle \right\} \right] \\ &\leq \mathbb{E}_\varepsilon \left[\sup_{\mathbf{w} \in \mathcal{W}} \left\{ \frac{2\sqrt{3H}}{m} \sum_{i=1}^m \varepsilon_i \sqrt{\tau_i} \langle \mathbf{w}, \mathbf{x}_i \rangle \right\} \right] + \mathbb{E}_\varepsilon \left[\sup_{\beta \in \mathcal{V}} \left\{ \frac{2\sqrt{3H}}{m} \sum_{i=1}^m \varepsilon_i \sqrt{\tau_i} \langle \beta, \mathbf{f}(\mathbf{x}_i) \rangle \right\} \right] \\ &= \hat{\mathfrak{R}}_{S'}(\mathcal{H}_\mathcal{W}) + \hat{\mathfrak{R}}_{S''}(\mathcal{H}_\mathcal{V}). \end{aligned}$$

Having this, we will follow the proof of Kakade et al. [2012, Theorem 7] to bound the empirical Rademacher complexities $\hat{\mathfrak{R}}_{S'}(\mathcal{H}_\mathcal{W})$ and $\hat{\mathfrak{R}}_{S''}(\mathcal{H}_\mathcal{V})$ with quantities of interest. Let $t > 0$ and apply Lemma 9 with $\mathbf{u} = \mathbf{w}$ and $\mathbf{v}_i = t\varepsilon_i\sqrt{\tau_i}\mathbf{x}_i$ to get

$$\begin{aligned} &\sup_{\mathbf{w} \in \mathcal{W}} \left\{ \sum_{i=1}^m \langle \mathbf{w}, t\varepsilon_i\sqrt{\tau_i}\mathbf{x}_i \rangle \right\} \\ &\leq \frac{t^2}{2\sigma} \sum_{i=1}^m \|\varepsilon_i\sqrt{\tau_i}\mathbf{x}_i\|_*^2 + \sup_{\mathbf{w} \in \mathcal{W}} \Omega(\mathbf{w}) + \sum_{i=1}^m \langle \nabla \Omega^*(\mathbf{v}_{1:i-1}), \varepsilon_i\sqrt{\tau_i}\mathbf{x}_i \rangle \\ &\leq \frac{t^2 B^2}{2\sigma} \sum_{i=1}^m |\tau_i| + \alpha \hat{R}_S(f) + \sum_{i=1}^m \langle \nabla \Omega^*(\mathbf{v}_{1:i-1}), \varepsilon_i\sqrt{\tau_i}\mathbf{x}_i \rangle. \end{aligned}$$

Now take expectation w.r.t. all the ε_i on both sides. The left hand side is $mt\hat{\mathfrak{R}}_{S'}(\mathcal{H}_\mathcal{W})$ and the last term on the right hand side becomes zero since $\mathbb{E}[\varepsilon_i] = 0$. Denoting $r = \frac{1}{m} \sum_{i=1}^m |\tau_i|$ and multiplying through by $\frac{1}{mt}$, we get

$$\hat{\mathfrak{R}}_{S'}(\mathcal{H}_\mathcal{W}) \leq \frac{B^2 r t}{2\sigma} + \frac{\alpha}{mt} \hat{R}_S(f_\beta).$$

Proving analogously for $\hat{\mathfrak{R}}_{S''}(\mathcal{H}_\mathcal{V})$, we get that

$$\hat{\mathfrak{R}}_S(\mathcal{L}) \leq 2\sqrt{3H} \left(\frac{(B^2 + C^2)rt}{\sigma} + \frac{\alpha \hat{R}_S(f_\beta) + \rho}{mt} \right).$$

Optimizing over t gives us

$$\hat{\mathfrak{R}}_S(\mathcal{L}) \leq 4\sqrt{3H}(B+C) \sqrt{\frac{r(\alpha \hat{R}_S(f_\beta) + \rho)}{m\sigma}}.$$

Now focus on the upper bound of r . First we obtain bounds on each τ_i . We start with the bound on the loss function, exploiting smoothness. Let $\ell(\langle \mathbf{w}, \mathbf{x} \rangle + f_\beta(\mathbf{x}), y) = \phi(\langle \mathbf{w}, \mathbf{x} \rangle + f_\beta(\mathbf{x}))$, where $\phi : \mathbb{R} \mapsto \mathbb{R}$ is an H -smooth function. From the definition of smoothness [Shalev-Shwartz and Ben-David, 2014, (12.5)], we have that for all \mathbf{w} and \mathbf{v}

$$\begin{aligned} & \phi(\langle \mathbf{w}, \mathbf{x} \rangle + f_\beta(\mathbf{x})) \\ & \leq \phi(\langle \mathbf{v}, \mathbf{x} \rangle + f_\beta(\mathbf{x})) + \phi'(\langle \mathbf{v}, \mathbf{x} \rangle + f_\beta(\mathbf{x})) \langle \mathbf{w} - \mathbf{v}, \mathbf{x} \rangle + \frac{H}{2} \langle \mathbf{w} - \mathbf{v}, \mathbf{x} \rangle^2 \\ & \leq \phi(\langle \mathbf{v}, \mathbf{x} \rangle + f_\beta(\mathbf{x})) + 2\sqrt{H\phi(\langle \mathbf{v}, \mathbf{x} \rangle + f_\beta(\mathbf{x}))} \|\mathbf{w} - \mathbf{v}\| \|\mathbf{x}\|_* + \frac{H}{2} \|\mathbf{w} - \mathbf{v}\|^2 \|\mathbf{x}\|_*^2. \end{aligned} \quad (17)$$

To obtain the last inequality we used the generalized Cauchy-Schwarz inequality and the fact that for an H -smooth non-negative function ϕ , we have that $|\phi'(t)| \leq \sqrt{4H\phi(t)}$, [Srebro et al., 2010, Lemma 2.1]. Now recall a property of a σ -strongly-convex function F , that holds for its minimizer \mathbf{v} and any \mathbf{w} [Shalev-Shwartz and Ben-David, 2014, Lemma 13.5],

$$\|\mathbf{w} - \mathbf{v}\|^2 \leq \frac{2}{\sigma} (F(\mathbf{w}) - F(\mathbf{v})).$$

Since inequality (17) holds for any \mathbf{v} , set $\mathbf{v} = \mathbf{0}$, which is also the minimizer of $\Omega(\cdot)$, apply aforementioned property to get

$$\begin{aligned} & \phi(\langle \mathbf{w}, \mathbf{x} \rangle + f_\beta(\mathbf{x})) \leq \phi(f_\beta(\mathbf{x})) + 2\sqrt{\frac{2H}{\sigma} \phi(f_\beta(\mathbf{x})) \Omega(\mathbf{w})} \|\mathbf{x}\|_* + \frac{H}{\sigma} \Omega(\mathbf{w}) \|\mathbf{x}\|_*^2 \\ \Rightarrow & \ell(\langle \mathbf{w}, \mathbf{x}_i \rangle + f_\beta(\mathbf{x}_i), y_i) \leq \tau_i \end{aligned} \quad (18)$$

$$= \ell(f_\beta(\mathbf{x}_i), y_i) + \sqrt{\frac{8HB^2\alpha}{\sigma} \hat{R}_S(f_\beta) \ell(f_\beta(\mathbf{x}_i), y_i)} + \frac{HB^2\alpha}{\sigma} \hat{R}_S(f_\beta). \quad (19)$$

The last inequality comes from the definition of the class \mathcal{H} . Now we consider the average and, by Jensen's inequality,

$$\begin{aligned} r &= \frac{1}{m} \sum_{i=1}^m |\tau_i| = \hat{R}_S(f_\beta) + \frac{1}{m} \sum_{i=1}^m \sqrt{\frac{8HB^2\alpha}{\sigma} \hat{R}_S(f_\beta) \ell(f_\beta(\mathbf{x}_i), y_i)} + \frac{HB^2\alpha}{\sigma} \hat{R}_S(f_\beta) \\ &\leq \hat{R}_S(f_\beta) + \sqrt{\frac{8HB^2\alpha}{\sigma} \hat{R}_S(f_\beta)} + \frac{HB^2\alpha}{\sigma} \hat{R}_S(f_\beta) \leq \left(1 + \sqrt{\frac{2HB^2\alpha}{\sigma}}\right)^2 \hat{R}_S(f_\beta). \end{aligned}$$

This gives us

$$\begin{aligned} \hat{\mathfrak{R}}_S(\mathcal{L}) &\leq 4\sqrt{3H}(B+C) \left(1 + \sqrt{\frac{2HB^2\alpha}{\sigma}}\right) \sqrt{\frac{\hat{R}_S(f_\beta)(\alpha \hat{R}_S(f_\beta) + \rho)}{m\sigma}} \\ &\leq 4\sqrt{3H}(B+C) \left(1 + \sqrt{\frac{2HB^2\alpha}{\sigma}}\right) \frac{\hat{R}_S(f_\beta)\sqrt{\alpha} + \sqrt{\hat{R}_S(f_\beta)\rho}}{\sqrt{m\sigma}}. \end{aligned}$$

Taking expectation w.r.t. the sample on both sides and applying Jensen's inequality gives the statement. ■

6.3 Proofs of Main Results

Proof of Theorem 2. To show the statement we will apply Theorem 7. In particular, we will consider any choice of \mathbf{w} and β within the set induced by a strongly-convex function Ω . To apply Theorem 7, we need to upper bound the Rademacher complexity of the loss class \mathcal{L} and also the quantity $r = \sup_{f \in \mathcal{L}} \mathbb{E}_{(\mathbf{x}, y)} [f(\mathbf{x}, y)]$.

We obtain the bound on Rademacher complexity by applying Theorem 10. First define the loss class $\mathcal{L} := \{(\mathbf{x}, y) \mapsto \ell(h, y) : h \in \mathcal{H}\}$, and hypothesis class

$$\mathcal{H} := \left\{ \mathbf{x} \mapsto \langle \mathbf{w}, \mathbf{x} \rangle + h_{\beta}^{\text{src}}(\mathbf{x}) : \right. \\ \left. \Omega(\mathbf{w}) \leq \frac{1}{\lambda} \hat{R}_S(h_{\beta}^{\text{src}}) \wedge \Omega(\beta) \leq \rho \wedge \hat{R}_S(h_{\mathbf{w}, \beta}) \leq \hat{R}_S(h_{\beta}^{\text{src}}) \right\}.$$

To motivate the choice for the constraints observe that for

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmin}} \left\{ \hat{R}_S(h_{\mathbf{w}, \beta}) + \lambda \Omega(\mathbf{w}) \right\},$$

we have $\Omega(\hat{\mathbf{w}}) \leq \lambda^{-1} \hat{R}_S(h_{\mathbf{0}, \beta}) = \lambda^{-1} \hat{R}_S(h_{\beta}^{\text{src}})$, and $\hat{R}_S(h_{\hat{\mathbf{w}}, \beta}) \leq \hat{R}_S(h_{\beta}^{\text{src}})$. That said, by applying Theorem 10 with $\alpha = \frac{1}{\lambda}$ and $f_{\beta} = h_{\beta}^{\text{src}}$ and assuming that $\lambda \leq \kappa$, we obtain

$$\mathfrak{R}(\mathcal{L}) \leq \mathcal{O} \left(\frac{R^{\text{src}} \kappa}{\sqrt{m\lambda}} + \sqrt{\frac{R^{\text{src}} \rho \kappa^2}{m\lambda}} \right).$$

Next we obtain the bound on r

$$r = \sup_{h \in \mathcal{H}} \mathbb{E}_{(\mathbf{x}, y)} [\ell(h(\mathbf{x}), y)] = \sup_{h \in \mathcal{H}} \mathbb{E}_S [\hat{R}_S(h)] \leq \mathbb{E}_S \left[\sup_{h \in \mathcal{H}} \hat{R}_S(h) \right] \leq \mathbb{E}_S [\hat{R}_S(h_{\beta}^{\text{src}})] = R^{\text{src}}.$$

The last two inequalities come from Jensen's inequality and the definition of the class \mathcal{H} . Plugging the bounds on the Rademacher complexity and r into the statement of Theorem 7, and applying the inequality $\sqrt{a+b} \leq \sqrt{a} + \frac{b}{2\sqrt{a}}$ to the \sqrt{v} term, gives the statement.

Proof of Theorem 3. For any choice of β with $\Omega(\beta) \leq \rho$, denote the best in the class by

$$\mathbf{w}^* = \underset{\mathbf{w} : \Omega(\mathbf{w}) \leq \tau}{\operatorname{argmin}} R(h_{\mathbf{w}, \beta}).$$

By the definition of $\hat{\mathbf{w}}$, we have

$$\hat{R}_S(h_{\hat{\mathbf{w}}, \beta}) + \lambda \Omega(\hat{\mathbf{w}}) \leq \hat{R}_S(h_{\mathbf{w}^*, \beta}) + \lambda \Omega(\mathbf{w}^*). \quad (20)$$

Now denote

$$Z = \kappa \sqrt{\frac{R^{\text{src}}}{m}} (\sqrt{R^{\text{src}}} + \sqrt{\rho}).$$

Then, by following the proof of Theorem 2 until the application of inequality $\sqrt{a+b} \leq \sqrt{a} + \frac{b}{2\sqrt{a}}$, ignoring constants, using the assumption (20), and assuming that $\lambda \leq \kappa \leq 1$ we have that

$$\begin{aligned} R(h_{\hat{\mathbf{w}}, \beta}) &\leq \hat{R}_S(h_{\mathbf{w}^*, \beta}) + \lambda \tau + \frac{Z}{\lambda} + \sqrt{\frac{M\eta}{m}} \sqrt{R^{\text{src}} + \frac{Z}{\lambda}} + \frac{M\eta}{m} \\ &\leq \hat{R}_S(h_{\mathbf{w}^*, \beta}) + \lambda \tau + \frac{Z}{\lambda} + \sqrt{\frac{R^{\text{src}} M\eta}{m}} + \frac{\sqrt{Z M\eta}}{\sqrt{m\lambda}} + \frac{M\eta}{m}. \end{aligned} \quad (21)$$

Optimizing the l.h.s. over λ gives

$$\lambda^* = \sqrt{\frac{Z}{\tau}} + \frac{1}{\tau} \sqrt{\frac{Z M\eta}{m}}.$$

We plug it back into (21) to obtain that

$$\begin{aligned}
R(h_{\mathbf{w}^*, \beta}) &\leq \hat{R}_S(h_{\mathbf{w}^*, \beta}) + \sqrt{\tau} \sqrt{Z + \sqrt{\frac{ZM\eta}{m}}} + \sqrt{\frac{R^{\text{src}}M\eta}{m}} + \frac{M\eta}{m} \\
&\leq \hat{R}_S(h_{\mathbf{w}^*, \beta}) + \frac{\sqrt{R^{\text{src}}} + \sqrt[4]{R^{\text{src}}\rho}}{\sqrt[4]{m}} \sqrt{\kappa\tau} + \frac{\sqrt[4]{R^{\text{src}}} + \sqrt[8]{R^{\text{src}}\rho}}{\sqrt[4]{m^{1.5}}} \sqrt[4]{\kappa\tau^2 M\eta} \\
&\quad + \sqrt{\frac{R^{\text{src}}M\eta}{m}} + \frac{M\eta}{m}.
\end{aligned} \tag{22}$$

All that is left is to concentrate $\hat{R}_S(h_{\mathbf{w}^*, \beta})$ around its mean. Denoting the variance by

$$V = \mathbb{E} \left[\sum_{i=1}^m (\ell(h_{\mathbf{w}^*, \beta}(\mathbf{x}_i), y_i) - R(h_{\mathbf{w}^*, \beta}))^2 \right],$$

we apply Bernstein's inequality

$$\mathbb{P} \left(\sum_{i=1}^m (\ell(h_{\mathbf{w}^*, \beta}(\mathbf{x}_i), y_i) - R(h_{\mathbf{w}^*, \beta})) > t \right) \leq \exp \left(-\frac{t^2/2}{V + Mt/3} \right).$$

Setting

$$e^{-\eta} = \exp \left(-\frac{t^2/2}{V + Mt/3} \right),$$

we have that with probability at least $1 - e^{-\eta}$, $\forall \eta \geq 0$

$$\begin{aligned}
\hat{R}_S(h_{\mathbf{w}^*, \beta}) &\leq R(h_{\mathbf{w}^*, \beta}) + \sqrt{\frac{2\eta \mathbb{E} [(\ell(h_{\mathbf{w}^*, \beta}(\mathbf{x}_i), y_i) - R(h_{\mathbf{w}^*, \beta}))^2]}{m}} + \frac{2M\eta}{3m} \\
&\leq R(h_{\mathbf{w}^*, \beta}) + 2\sqrt{\frac{R(h_{\mathbf{w}^*, \beta})M\eta}{m}} + \frac{2M\eta}{3m} \\
&\leq R(h_{\mathbf{w}^*, \beta}) + 2\sqrt{\frac{R^{\text{src}}M\eta}{m}} + \frac{2M\eta}{3m}.
\end{aligned}$$

The last inequality comes from the observation that $R(h_{\mathbf{w}^*, \beta}) \leq R(h_0) = R^{\text{src}}$. Plugging this result into (22) completes the proof.

7 Conclusions

In this paper we have formally captured and theoretically analyzed a general family of learning algorithms transferring information from multiple supplied source hypotheses. In particular, our formulation stems from the regularized Empirical Risk Minimization principle with the choice of any non-negative smooth loss function and any strongly convex regularizer. Theoretically we have analyzed the generalization ability and excess risk of this family of HTL algorithms. Our analysis showed that a good source hypothesis combination facilitates faster generalization, specifically in $\mathcal{O}(1/m)$ instead of the usual $\mathcal{O}(1/\sqrt{m})$. Furthermore, given a perfect source hypothesis combination, our analysis is consistent with the intuition that learning is not required. As a byproduct of our investigation, we came up with new results in Rademacher complexity analysis of the smooth loss classes, that could be of independent interest.

Our conclusions suggest the key importance of a source hypothesis selection procedure. Indeed, when an algorithm is provided with enormous pool of source hypotheses, how to select relevant ones on the basis of only few labeled examples? This might sound similar to the feature selection problem under the condition that $n \gg m$, however, earlier empirical studies by Tommasi et al. [2014] with hundreds of sources did not find much corroboration for this hypothesis when applying $L1$ regularization. Thus, it remains unclear if having few good sources from hundreds is a reasonable assumption.

Acknowledgements

This work is partially supported by the Swiss National Science Foundation Sinergia project Ninapro (I.K.).

References

- Y. Aytar and A. Zisserman. Tabula rasa: Model transfer for object category detection. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 2252–2259. IEEE, 2011.
- P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2003.
- P. L. Bartlett, O. Bousquet, and S. Mendelson. Local rademacher complexities. *Annals of Statistics*, pages 1497–1537, 2005.
- S. Ben-David and R. Urner. On the hardness of domain adaptation and the utility of unlabeled target samples. In *Algorithmic Learning Theory*, pages 139–153. Springer, 2012.
- S. Ben-David and R. Urner. Domain adaptation as learning with auxiliary information. In *New Directions in Transfer and Multi-Task - Workshop @ Advances in Neural Information Processing Systems*, 2013.
- S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- A. Bergamo and L. Torresani. Classes and other classifier-based features for efficient object categorization. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PP(99), 2014.
- A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. In *Conference on Computational learning theory*, pages 92–100. ACM, 1998.
- O. Bousquet. *Concentration inequalities and empirical processes theory applied to the analysis of learning algorithms*. PhD thesis, Ecole Polytechnique, 2002.
- O. Bousquet and A. Elisseeff. Stability and Generalization. *Journal of Machine Learning Research*, 2: 499–526, 2002.
- R. Caruana. *Multitask learning*. Springer, 1998.
- O. Chapelle, B. Schölkopf, A. Zien, et al. *Semi-supervised learning*, volume 2. MIT press Cambridge, 2006.
- C. Cortes and M. Mohri. Domain adaptation and sample bias correction theory and algorithm for regression. *Theoretical Computer Science*, 519:103–126, 2014.
- H. Daumé III. Frustratingly easy domain adaptation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics*, 2007.
- H. Daumé III, A. Kumar, and A. Saha. Frustratingly easy semi-supervised domain adaptation. In *Proceedings of the 2010 Workshop on Domain Adaptation for Natural Language Processing*, pages 53–59. Association for Computational Linguistics, 2010.
- L. Duan, I. W. Tsang, D. Xu, and T.-S. Chua. Domain adaptation from multiple sources via auxiliary classifiers. In *International Conference on Machine Learning*, pages 289–296, 2009.
- A. Hoorfar and M. Hassani. Inequalities on the lambert w function and hyperpower function. *J. Inequal. Pure and Appl. Math*, 9(2), 2008.

- S. M. Kakade, K. Sridharan, and A. Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Advances in Neural Information Processing Systems 21*, pages 793–800, 2008.
- S. M. Kakade, S. Shalev-Shwartz, and A. Tewari. Regularization techniques for learning with matrices. *Journal of Machine Learning Research*, 13:1865–1890, 2012.
- W. Kienzle and K. Chellapilla. Personalized handwriting recognition via biased regularization. In *International Conference on Machine Learning*, pages 457–464, 2006.
- I. Kuzborskij and F. Orabona. Stability and Hypothesis Transfer Learning. In *International Conference on Machine Learning*, pages 942–950, 2013.
- I. Kuzborskij, F. Orabona, and B. Caputo. From N to N+1: multiclass transfer incremental learning. In *Conference on Computer Vision and Pattern Recognition*, pages 3358–3365, 2013.
- I. Kuzborskij, F. Orabona, and B. Caputo. Transfer Learning through Greedy Subset Selection. In *International Conference on Image Analysis and Processing*, 2015.
- L.-J. Li, H. Su, E. P. Xing, and L. Fei-Fei. Object bank: A high-level image representation for scene classification & semantic feature sparsification. In *Advances in Neural Information Processing Systems 23*, pages 1378–1386, 2010.
- X. Li and J. Bilmes. A bayesian divergence prior for classifier adaptation. In *International Conference on Artificial Intelligence and Statistics*, pages 275–282, 2007.
- J. Luo, T. Tommasi, and B. Caputo. Multiclass transfer learning from unconstrained priors. In *International Conference on Computer Vision*, pages 1863–1870, 2011.
- Y. Mansour, M. Mohri, and A. Rostamizadeh. Domain adaptation with multiple sources. In *Advances in Neural Information Processing Systems 21*, pages 1041–1048, 2008.
- Y. Mansour, M. Mohri, and A. Rostamizadeh. Domain adaptation: Learning bounds and algorithms. In *The Conference on Learning Theory*, 2009.
- M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of machine learning*. The MIT Press, 2012.
- M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *Conference on Computer Vision and Pattern Recognition*, pages 1717–1724, 2014.
- F. Orabona, C. Castellini, B. Caputo, A.E. Fiorilla, and G. Sandini. Model Adaptation with Least-Squares SVM for Adaptive Hand Prosthetics. In *Robotics and Automation, IEEE International Conference on*, pages 2897–2903. IEEE, 2009.
- S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.*, 22(10):1345–1359, 2010.
- B. Schölkopf, R. Herbrich, and A. J. Smola. A generalized representer theorem. In *Conference on Computational learning theory*, pages 416–426. Springer, 2001.
- S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- V. Sharmanska, N. Quadrianto, and C. H. Lampert. Learning to rank using privileged information. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 825–832. IEEE, 2013.

- N. Srebro, K. Sridharan, and A. Tewari. Smoothness, low noise and fast rates. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 2199–2207, 2010.
- M. E. Taylor and P. Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10:1633–1685, 2009.
- S. Thrun and L. Pratt. *Learning to Learn*. Springer Science & Business Media, 1998.
- T. Tommasi, F. Orabona, and B. Caputo. Safety in numbers: Learning categories from few examples with multi model knowledge transfer. In *The Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2010, San Francisco, CA, USA, 13-18 June 2010*, pages 3081–3088, 2010.
- T. Tommasi, F. Orabona, C. Castellini, and B. Caputo. Improving control of dexterous hand prostheses using adaptive learning. *Robotics, IEEE Transactions on*, 29(1):207–219, 2013.
- T. Tommasi, F. Orabona, and B. Caputo. Learning categories from few examples with multi model knowledge transfer. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(5):928–941, 2014.
- V. Vapnik and A. Vashist. A new learning paradigm: Learning using privileged information. *Neural Networks*, 22(5):544–557, 2009.
- E. De Vito, A. Caponnetto, and L. Rosasco. Model selection for regularized least-squares algorithm in learning theory. *Foundations of Computational Mathematics*, 5(1):59–85, 2005.
- J. Yang, R. Yan, and A.G. Hauptmann. Cross-Domain Video Concept Detection Using Adaptive SVMs. In *Proceedings of the 15th international conference on Multimedia*, pages 188–197. ACM, 2007.

A Additional Proofs

Theorem 11 Let $h_{\hat{\mathbf{w}},\beta}$ be generated by Regularized ERM, given the m -sized training set S sampled i.i.d. from the target domain, source hypotheses $\{h_i^{\text{src}}\}_{i=1}^n$, any source weights β obeying $\Omega(\beta) \leq \rho$, and $\lambda \in \mathbb{R}_+$. Assume that $\ell(h_{\hat{\mathbf{w}},\beta}(\mathbf{x}), y) \leq M$ for any (\mathbf{x}, y) and any training set. Then, denoting $\kappa = \frac{H}{\sigma}$ and assuming that $\lambda \leq 1$, we have with probability at least $1 - e^{-\eta}$, $\forall \eta \geq 0$

$$R(h_{\hat{\mathbf{w}},\beta}) \leq \hat{R}_S(h_{\hat{\mathbf{w}},\beta}) + \tilde{\mathcal{O}}\left(\left(\sqrt{\frac{R^{\text{src}}}{m}} + \sqrt[4]{\frac{M^2\rho}{m^3\sigma}} + \sqrt[8]{\frac{M^4\rho}{m^7\lambda^2\sigma^3}}\right)\left(\sqrt{\frac{M\kappa}{\lambda}} + \sqrt{\kappa\rho} + \sqrt{M\eta}\right)\right).$$

Proof To prove the statement we will use Theorem 1 of Srebro et al. [2010]. In particular, we need to obtain bounds on the empirical risk and also to bound the worst case Rademacher complexity of the class

$$\mathcal{H} = \left\{ \mathbf{x} \mapsto \langle \mathbf{w}, \mathbf{x} \rangle + h_{\beta}^{\text{src}}(\mathbf{x}) : \Omega(\mathbf{w}) \leq \frac{\hat{R}_S(h_{\beta}^{\text{src}})}{\lambda} \wedge \Omega(\beta) \leq \tau \right\}.$$

The corresponding loss class is

$$\mathcal{L} = \{(\mathbf{x}, y) \mapsto \ell(h(\mathbf{x}), y) : h \in \mathcal{H} \wedge R(h) \leq R^{\text{src}}\}.$$

A constraint on $\Omega(\beta)$ in \mathcal{H} comes from the statement of the theorem, while a constraint on $\Omega(\hat{\mathbf{w}})$ comes from an observation that for

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmin}} \left\{ \hat{R}_S(h_{\mathbf{w},\beta}) + \lambda\Omega(\mathbf{w}) \right\},$$

so we have $\Omega(\hat{\mathbf{w}}) \leq \frac{\hat{R}_S(h_{\mathbf{0},\beta})}{\lambda}$. The same argument immediately gives us a bound on the empirical risk, that is, $\hat{R}_S(h_{\hat{\mathbf{w}},\beta}) \leq \hat{R}_S(h_{\mathbf{0},\beta}) = \hat{R}_S(h_{\beta}^{\text{src}})$. Taking expectation on both sides gives the constraint of \mathcal{L} .

By applying Theorem 1 of Kakade et al. [2008] and subadditive property of Rademacher complexities [Bartlett and Mendelson, 2003], we have that

$$\hat{\mathfrak{R}}_S(\mathcal{H}) \leq \sqrt{\frac{2\hat{R}_S(h_{\beta}^{\text{src}})}{m\lambda\sigma}} + \sqrt{\frac{2\rho}{m\sigma}} \leq \sqrt{\frac{2M}{m\lambda\sigma}} + \sqrt{\frac{2\rho}{m\sigma}}. \quad (23)$$

Note that the upper bound is the bound on the worst-case Rademacher complexity since no term depends on the sample.

All that is left to do is to show the bound on the empirical risk in terms of R^{src} . However, we cannot use Theorem 1 of Srebro et al. [2010] since it is not symmetric. Instead we will use a similar localized bound of Bartlett et al. [2005, Corollary 3.5]. In order to apply it, we have to obtain an upper bound on the Rademacher complexity of the loss class \mathcal{L} that is a sub-root function [Bousquet, 2002, Definition 4.1]. By using the fact that loss function is bounded, we apply Talagrand's lemma [Mohri et al., 2012], have $\hat{\mathfrak{R}}_S(\mathcal{L}) \leq M\hat{\mathfrak{R}}_S(\mathcal{H})$, upper-bound with the first inequality of (23) and applying Jensen's inequality w.r.t. $\mathbb{E}[\cdot]$ have

$$\mathfrak{R}(\mathcal{L}) \leq M\sqrt{\frac{2R^{\text{src}}}{m\lambda\sigma}} + M\sqrt{\frac{2\rho}{m\sigma}}.$$

Since upper bound is a sub-root function of R^{src} , we obtain it's fixed point r^* as required by Corollary 3.5 and conclude that

$$r^* \leq \sqrt{\frac{2M^2\rho}{m\sigma}} + \frac{2M^2}{m\lambda\sigma} + 2M\sqrt[4]{\frac{8\rho}{m^3\lambda^2\sigma^3}}.$$

Now we apply Corollary 3.5 and for any $K > 0$ we have with probability at least $1 - e^{-\eta}$, $\forall \eta \geq 0$ the following holds

$$\hat{R}_S(h_{\hat{\mathbf{w}},\beta}) \leq K \left(R^{\text{src}} + \sqrt{\frac{M^2\rho}{m\sigma}} + \frac{M^2}{m\lambda\sigma} + M\sqrt[4]{\frac{\rho}{m^3\lambda^2\sigma^3}} + \frac{1+\eta}{m} \right).$$

All that is left to do is to apply Theorem 1 of Srebro et al. [2010] to have

$$\begin{aligned} R(h_{\hat{\mathbf{w}},\beta}) &\leq \hat{R}_S(h_{\hat{\mathbf{w}},\beta}) + \tilde{O} \left(\left(\sqrt{\frac{R^{\text{src}}}{m}} + \sqrt[4]{\frac{M^2\rho}{m^3\sigma}} + \frac{M}{m\sqrt{\lambda\sigma}} + \sqrt[8]{\frac{M^2\rho}{m^7\lambda^2\sigma^3}} + \frac{\sqrt{1+\eta}}{m} \right) \times \right. \\ &\quad \left. \times \left(\sqrt{\frac{M\kappa}{\lambda}} + \sqrt{\kappa\rho} + \sqrt{M\eta} \right) + \frac{M\kappa}{m\lambda} + \frac{\kappa\rho}{m} \right). \end{aligned}$$

Using the assumption on λ , we get the stated result. ■

A.1 Guarantees using Localized Rademacher Complexity Bounds

The following theorem is due to Bousquet [2002, Theorem 6.1]. In particular, we state the inequality appearing prior to the last in the proof, as it better serves our purpose.

Theorem 12 (Bousquet [2002]) *Let \mathcal{F} be a class of non-negative functions such that $\|f\|_{\infty} \leq M$ almost surely. Let ϕ_m be a function defined on $[0, \infty)$ that is non-negative, non-decreasing, not identically zero, and such that $\phi_m(r)/\sqrt{r}$ is non-increasing. Moreover let ϕ_m be such that for all $r > 0$*

$$\hat{\mathfrak{R}}_S(\mathcal{F}) \leq \phi_m(r).$$

Define r_m^* as the largest solution of the equation $\phi_m(r) = r$. Then, for all $\eta > 0$, with probability at least $1 - e^{-\eta}$ for all $f \in \mathcal{F}$ and any $\{X_i\}_{i=1}^m$ drawn i.i.d.

$$\begin{aligned} \mathbb{E}_X[f(X)] &\leq \frac{1}{m} \sum_{i=1}^m f(X_i) + 45r_m^* + \sqrt{8r_m^* \mathbb{E}_X[f(X)]} + \sqrt{\frac{4M(\eta + 6 \log \log m) \mathbb{E}_X[f(X)]}{m}} \\ &\quad + \frac{20M(\eta + 6 \log \log m)}{m}. \end{aligned}$$

The following HTL generalization bound is shown using Theorem 12.

Theorem 13 Let $h_{\mathbf{w}, \beta}$ be generated by Regularized ERM, given the m -sized training set S sampled i.i.d. from the target domain, source hypotheses $\{h_i^{\text{src}}\}_{i=1}^n$, any source weights β obeying $\Omega(\beta) \leq \rho$, and $\lambda \in \mathbb{R}_+$. Assume that ℓ is a L -Lipschitz loss function and $\ell(h_{\mathbf{w}, \beta}(\mathbf{x}), y) \leq M$ for any (\mathbf{x}, y) and any training set. Then we have with probability at least $1 - e^{-\eta}$, $\forall \eta \geq 0$

$$\begin{aligned} R(h_{\mathbf{w}, \beta}) &\leq \hat{R}_S(h_{\mathbf{w}, \beta}) + \tilde{O} \left(\frac{L^2 + L}{m\lambda\sigma} + L\sqrt{\frac{\rho}{m\sigma}} + \sqrt{\frac{R^{\text{src}}(L^2 + L)}{m\lambda\sigma}} \right. \\ &\quad \left. + \sqrt{R^{\text{src}}} \sqrt{\frac{L^2\rho}{m\sigma}} + \sqrt{\frac{R^{\text{src}}M\eta}{m}} + \frac{M\eta}{m} \right). \end{aligned}$$

Proof The core of the proof is an application of Theorem 12. In particular, we have to obtain the fixed point r_m^* and upper bound $R(h)$ with the risk of the source hypothesis R^{src} .

Considering the L -Lipschitz loss class of Theorem 12 to be $\mathcal{L} := \{(\mathbf{x}, y) \mapsto \ell(h(\mathbf{x}), y) : h \in \mathcal{H}\}$, we have the relationship $\hat{\mathfrak{R}}_S(\mathcal{L}) \leq L\hat{\mathfrak{R}}_S(\mathcal{H})$ via Talagrand's lemma [Mohri et al., 2012, Lemma 4.2]. Furthermore, let the hypothesis class be

$$\mathcal{H} = \left\{ \mathbf{x} \mapsto \langle \mathbf{w}, \mathbf{x} \rangle + h_{\beta}^{\text{src}}(\mathbf{x}) : \Omega(\mathbf{w}) \leq \frac{1}{\lambda} \hat{R}_S(h_{\beta}^{\text{src}}) \wedge \Omega(\beta) \leq \rho \wedge \hat{R}_S(h_{\mathbf{w}, \beta}) \leq \hat{R}_S(h_{\beta}^{\text{src}}) \right\}.$$

The motivation for the choice of constraints comes from the same argument as in the proof of Theorem 2. That said, we obtain the upper bound

$$\hat{\mathfrak{R}}_S(\mathcal{L}) \leq L\sqrt{\frac{2R^{\text{src}}}{m\lambda\sigma}} + L\sqrt{\frac{2\rho}{m\sigma}}.$$

Both terms come by applying Theorem 7 by Kakade et al. [2012]. In the first term we set $f_{\max} = R^{\text{src}}$ and in the second $f_{\max} = \rho$. Now define function $\phi_m(r) = L\sqrt{\frac{2r}{m\lambda\sigma}} + L\sqrt{\frac{2\rho}{m\sigma}}$, and observe that it verifies the condition of Theorem 12. Next, to obtain the upper bound on r_m^* , we solve $L\sqrt{\frac{2r}{m\lambda\sigma}} + \frac{2\rho}{m\sigma} \leq r$ and get that $r_m^* \leq \frac{L(L+1)}{m\lambda\sigma} + L\sqrt{\frac{2\rho}{m\sigma}}$. As in Theorem 2, we also get that $R(h) \leq R^{\text{src}}$. Plugging r_m^* and the bound on $R(h)$ into Theorem 12, we have the statement. \blacksquare